34th Annual International Conference of the IEEE EMBS
San Diego, California USA, 28 August - 1 September, 2012

Dual-Modal Visibility Metrics for Interactive PET-CT Visualization

Younhyun Jung, Jinman Kim, *Member, IEEE*, and David Dagan Feng, *Fellow, IEEE*

***Abstract*— Dual-modal positron emission tomography and computed tomography (PET-CT) imaging enables the visualization of functional structures (PET) within human bodies in the spatial context of their anatomical (CT) counterparts, and is providing unprecedented capabilities in understanding diseases. However, the need to access and assimilate the two volumes simultaneously has raised new visualization challenges. In typical dual-modal visualization, the transfer functions for the two volumes are designed in isolation with the resulting volumes being fused. Unfortunately, such transfer function design fails to exploit the correlation that exists between the two volumes. In this study, we propose a dual-modal visualization method where we employ 'visibility' metrics to provide interactive visual feedback regarding the occlusion caused by the first volume on the second volume and vice versa. We further introduce a region of interest (ROI) function that allows visibility analyses to be restricted to subsection of the volume. We demonstrate the new visualization enabled by our proposed dual-modal visibility metrics using clinical whole-body PET-CT studies of various diseases.**

I. Introduction

The next-generation of dual-modality images, such as the co-aligned positron emission tomography and computed tomography (PET-CT) scans, have introduced new and improved capabilities in cancer staging and assessing responses to treatment [1, 2]. PET-CT images enable the visualization of functional tumor activity (PET) in the context of its surrounding anatomical structures (CT). However, accessing and assimilating critical information within these dual-modal volumes is becoming difficult with conventional 2D and even with 3D visualization approaches due to the massive amount of images acquired during PET-CT study.

The most commonly used 3D visualization is direct volume rendering, where every sample in a volume is mapped to the optical properties of opacity and color via a transfer function [3, 4]. Transfer functions are a key factor in designing the visual quality of images generated via direct volume rendering. There has been extensive research with the aim to improve the designing of the transfer function via various quality metrics e.g. first or second order derivatives [5, 6], use of size and occlusion attributes [7, 8], and texture feature based transfer function [9].

However, these algorithms are inherently optimized for single-modal volumes, and their applications to dual-modal volumes are often restricted to the design of the transfer function for individual volumes. One of the first dual-modal rendering was presented by Cai and Sakas [10] who presented a data intermixing algorithm to fuse two volumes in the rendering process. In a recent study by Bramon et al. [11], an information-theoretic approach that automatically selected the most informative voxels from two volumes was proposed and evaluated using medical data sets. The study reported promising results during fusion but only with the voxel's intensity and without its opacity, thus limiting the visualization to cross-sectional 2D views. In another study, Kim et al. [12] presented a dual-modal transfer function in the fused visualization of PET-CT volumes. In that study, a pair of 1D transfer functions was used to individually set the color and opacity of the PET and CT volumes where the combined rendering was able to visualize the ROI from PET in the context of its anatomical counterpart from CT. In these studies, transfer functions were applied to individual volumes or to fused volumes, and therefore disregarded how a change in the transfer function of one volume may affect the second volume. Our motivation stems from the lack of feedback in the design of dual-modal transfer functions.

We propose a new algorithm to equip the user with a feedback mechanism for the transfer function design process. We achieve this by providing the user with the visibility attribute that measures the level of occlusion of a sample in the volume. The visibility of a sample is a fundamental optical property that has been used in numerous visualization applications e.g. in the construction of a occlusion spectrum to improve the design of 2D transfer function in volume rendering [8], to guide the selection of the optimal view-point of a volume [13], and controlling the visibility to enhance the ROIs in a volume via non-photorealistic rendering [14]. In a recent study, Carlos and Ma [15] introduced a novel application of visibility to build a visibility histogram, which is a multi-dimensional graphical representation of the distribution of visibility of the samples in the volume rendered images. The histogram was used to guide the design of the transfer function for several single-modal volumes. We adopt the visibility to indicate the level of occlusion from dual-modal images to assist the transfer function design process.

II. Dual-Modal Visibility Histogram (DM-VH)

A. Single-Modal Visibility Histogram

The visibility of a sample indicates the contribution of that sample to a rendered volume with regards to its opacity [15]. This visibility $T(p)$ is the accumulation of opacities from an eye position E to a sample position p according to:

$$T(p) = e^{-\int_p^E O(t)dt} \quad (1)$$

where $O(t)$ is the opacity transfer function assigned to a sample. In this context, the visibility of a sample depends on the user-defined transfer function. The calculated visibilities of the samples are then added into its corresponding intensity bins of the visibility histogram:

This research was funded in part by ARC grants.

Younhyun Jung, Jinman Kim and David Dagan Feng are with the School of Information Technologies, University of Sydney, Australia (e-mail: yjun6175@uni.sydney.edu.au), Jinman Kim (phone: +61 2 9036 9708; fax: +61 2 9351 3838; e-mail: jinman.kim@sydney.edu.au).

David Dagan Feng is also with the CMSP, Department of EIE, Hong Kong Polytechnic University, Hong Kong, and Med-X Research Institute, Shanghai Jiao Tong University, China (e-mail: david.feng@sydney.edu.au).

$$VH[x] = VH[x] + O(p) \cdot T(p) \quad (2)$$

where *VH[x]* is the x^{th} bin of the visibility histogram, $O(p)$ is an opacity of a sample in position p, and $T(p)$ is the visibility of the sample. A visibility histogram thus represents the graphical distribution of the visibility in relation to the intensity bins of the histogram. Because the visibility histogram depends on view-point and opacity specification of transfer function, it is recalculated with every change in these properties.

B. Dual-Modal Visibility Histogram (dm-vh)

In the calculation of the dm-vh, the two volumes V_1 and V_2 are assumed to be aligned to the same coordinate space such that there are two sample values for every position p, respectively from V_1 and V_2. Using these sample pair, the combined visibility $T(p)$ of the two volumes are the result from the accumulation of opacity contribution, where 1.0 is the maximum opacity, from the two volumes calculated as:

$$T(p) = T(p) + \big(1 - T(p)\big) \cdot O(p)\,. \quad (3)$$

Fig. 1 illustrates the ray passing through the sample in V_1 followed by the sample in V_2 for each p, in a front-to-back composition. The resulting $O(p)$ is then used to build the visibility histogram using Eq. (2) for the two volumes.

C. DM-VH and Transfer Functions

Fig. 2 illustrates the effect of view dependency of dm-vh calculation, showing the correspondence between the rendered images and visibility values based on a view-point. Two aligned volumes A and B are composited and rendered. In (a), only the red sphere is visible as the opacity of the blue sphere is set to zero. The increase in the opacity of the blue sphere results in its visibility increasing while lowering the visibility of the red sphere (maximum visibility for the red sphere decreased from 15.52 to 13.17 (15.14%)), as shown in (b), effecting both its visual rendering and its visibility histograms. In (c), the volume is rotated such that the red sphere is now closer to the view-point and completely in front of the blue sphere, thus resulting in zero visibility for the blue sphere even though its opacity is set the same as in (b), while we see a large increase in the red sphere's visibility (150.64% increase from (a)). Finally, (d) is the opposite of (c) where the blue sphere is in front of the red sphere.

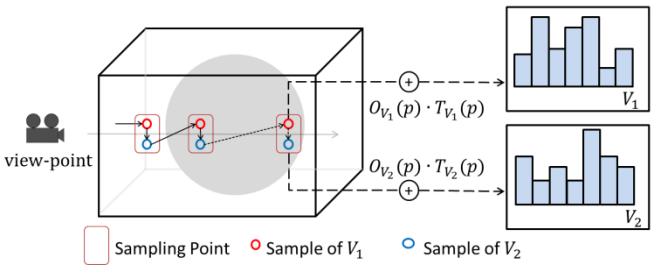

Figure 1. Computation of a dual-modal visibility histograms. Given a view-point and a position p, the accumulated opacity $T_V(p)$ from the two volumes V_1 and V_2 and the opacity of the current voxel O_V is multiplied with the resultant visibility being added to the corresponding bin of its volume's visibility histograms.

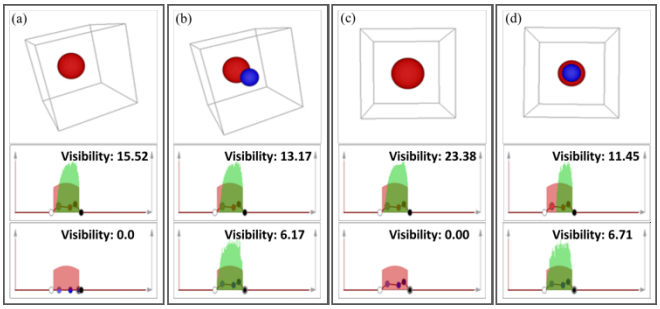

Figure 2. *Illustration of view dependency of a dual-modal visibility histogram on two aligned volumes A (larger red sphere; top visibility histogram) and B (smaller blue sphere occupying the same intensity distributions; bottom visibility histogram)*

D. Region of Interest (ROI) Subsection

The main drawback of the visibility histogram is that it is unable to differentiate and isolate ROIs from other non-related regions of the body with the same or similar intensity ranges. Without isolation, the visibility histogram fails to provide meaningful feedback in transfer function design. Fig. 3a is a rendering of PET-CT with its transfer function set for lung cancer: the CT is set to render the bones, skin, the lung tree and the boundary, all in low opacities; PET is set to render the high intensity ranges that are comprised of tumors but also other non-relevant structures including bladder, kidney, ureters below the lungs and brains above the lungs. The PET visibility histogram indicates high visibility for the tumor intensity range; this is due however, in part, to the visibility of other high intensity structures and not only of the lung. In fact, with this transfer function, it is difficult to see the large tumor residing in the left lung, which is clearly evident when the transfer function is modified in (c) and (d). In order to separate the ROI from non-relevant regions, we introduce a volume subsection function, which allows the user to interactively define a ROI within the volume.

1) Volume Clipping

Volume clipping plays a decisive role in understanding 3D volume because it allows interactive "cutting" away of parts of the volume that occludes volume that is of interest. Fig. 3(b) is an example of volume clipping being used to render only a part of the volume. Unlike the whole-body in (a), in this case the PET visibility histogram now only represents the lung and its tumors. The resulting visibility histogram can then be used to further refine the transfer functions.

2) Cuboid Selection

A cuboid can be interactively defined by the user as the ROI within the volume with the visibility histogram calculated only within the cuboid, as shown in Fig. 3(c). This capability ensures that the visibility histograms can be used to guide the transfer function to visualize the ROI in the context of seeing how the changes affect the full volume rendering. Fig. 3(d) shows the revised whole-body rendering using the transfer function in (c) to prioritize the visibility of the lung tumors.

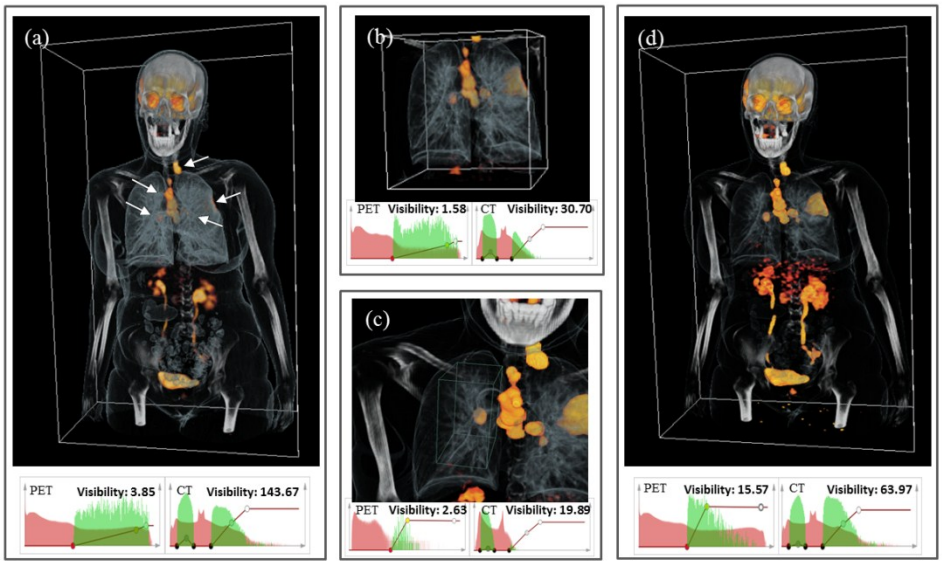

Figure 3. *(a) Whole-body fused PET-CT study diagnosed with lung cancer spread around both the left and right lungs, the bronchus, and also on the neck (arrows). The ROI subsection using volume clipping (b) isolates the lung regions from other non-relevant regions, thus resulting in PET's visibility histogram that only highlights the tumors. The cuboid in (c) is an example of how an ROI can be selected within the context of its whole-body rendering, which can then be used to guide the design of the transfer functions to refine the visualization of the tumors in (d).*

E. DM-VH implementation

1) Visibility Histogram Computation

As a visual feedback mechanism, dm-vh needs to be computed in interactive volume rendering. To achieve real-time performance, we used the GPU computing [16]. We used the Voreen volume rendering library [17], an open source engine that allows interactive visualization of volumetric data sets with high flexibility for integrating new visualization techniques. Texture-based volume rendering was used to render the dual-modal volumes using view-aligned slices from a given view-point. In the rendering pipeline, we used a fragment shader to calculate the visibility of the samples from the two volumes with a single raycast. Parallel processing, together with the large capacity in vertex fragment streams were used to calculate the visibility of several slices simultaneously in a single-pass using frame buffer objects.

2) Computation Time

The computations were performed on a consumer PC equipped with an nVIDIA 8800 GTX graphics card with 768 MB, Intel Quad CPU@2.39GHz, 4.0GB RAM, and Windows 7 64-bit. We measured the dm-vh performance based on PET-CT volumes of 256x256x256. The frames per second (fps) were measured by at 3.1, averaged among the time taken to render the volumes (and the calculation of the visibility histogram pairs) during view-point rotations and transfer function manipulations.

III. Results and Discussions

A. PET-CT data pre-processing

All the PET-CT studies were conducted at a Siemens Biograph TruePoint PET-CT scanner in the department of PET and Nuclear Medicine, Royal Prince Alfred hospital. All the studies have 326 slices with the slice thickness of 3 mm to cover whole body from the top head to upper thigh with PET slice size of 168×168 at the pixel size of 4.07 mm^2 and CT slice size of 512×512 at the pixel size of 0.98 mm^2. The PET was resampled to the CT's dimensions. The CT scans were processed to remove the background and the scanner bed/linen automatically via adaptive thresholding and image subtraction from a bed template. The voxel intensity of the CT was in Hounsfield units (HU), which is a normalized intensity scale. For PET, we normalized intensity of PET with the standard uptake values (SUV) [18].

B. Whole-body PET-CT

Fig. 3 presents the dm-vh in the visualization of a whole-body PET-CT lung cancer study. Fig. 4 illustrates the dual-modality rendering of PET only (a), CT only (b) and their fusion (c). Without the dm-vh, the transfer function of the fused rendering is the same pair as in (a) and (b), therefore not providing any feedbacks in regards to the visibility of either PET or CT volumes. With the use of the visibility histogram, we have a visual indication of what parts of the PET and CT are visible as shown in the transfer functions in (c). Together with the ROI subsection function, reaching an ideal transfer function is more intuitive and efficient as shown in (d) and for the whole-body rendering in (e). In this example of a Lymphoma study, all the abnormalities were localized at the base of the back and in the axilla, meaning that the abnormalities are outside the lung organ and not in the bones. Using the PET's visibility histogram, the CT's transfer function was redesigned to visualize greater amount of the anatomy while retaining almost the same visibility of the abnormalities. The ability of our CT transfer function to render better anatomical definitions, as seen in other studies [7, 15], was restricted due to our PET-CT consisting of low-contrast CT volume.

In this study, we developed a ROI subsection function involving volume clipping and a cuboid selection. However, such manual approaches were limited due to the difficulty in accurately defining the ROI.

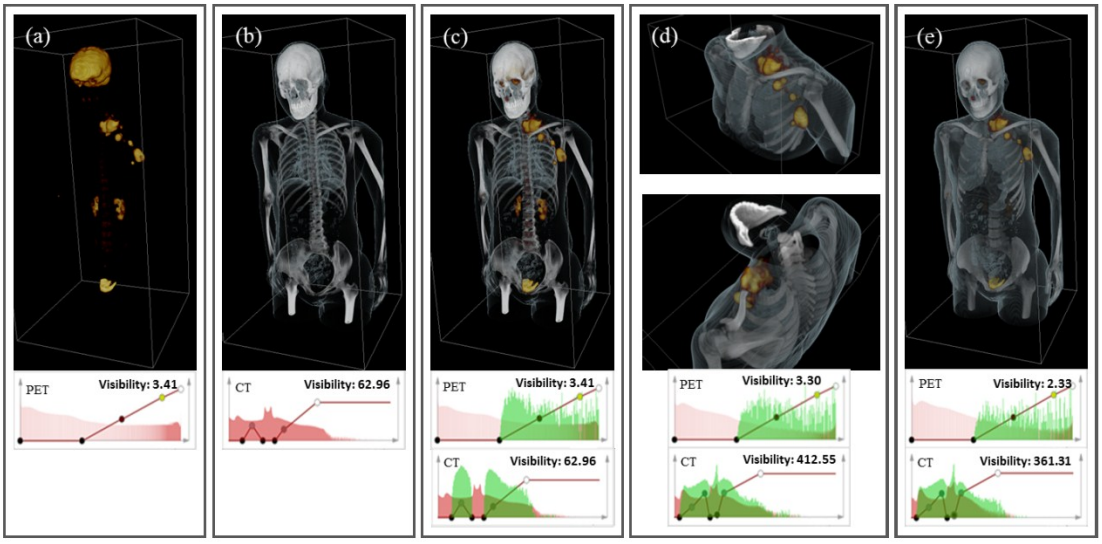

Figure 4. *Comparison of the visualization design using the conventional transfer function and with the addition of the visibility histogram.*

Weber et al. [19] presented a framework for using segmented regions in the transfer function design such that the opacity and color could be applied to only those voxels residing in the segments. We will investigate the use of segmentation results into the volume rendering pipeline to improve the ROI selection capability. The visibility histogram was used by Correa and Ma [15] to automatically optimize the transfer function design. The principle behind this was to iteratively find an optimal combination of the opacities among all the samples in the volume. We will investigate the application of visibility optimization together with our ROI function in our future studies. We anticipate that the visibility optimization from the ROI can lead to improved visualization that simplifies and improves the transfer function design process. Furthermore, we will harness the recent researches into transfer function automations as in [20] in our dual-modal visualization.

IV. Conclusions and Future Work

This paper described a new transfer function design mechanism for dual-modal PET-CT images. Through visual feedback during interactive volume rendering, the design of the transfer functions was enhanced and resulted in an improved depiction of the volume. Although only PET-CT images were used in this study, there are no restrictions in the modalities that can be used, and our future work will investigate the application to other multi-modal images including the emerging PET-MR images.

Acknowledgment

We like to thank our collaborators at the Royal Prince Alfred (RPA) hospital.

Reference

[1] M. Hutchings and S.F. Barrington, "PET/CT for therapy response assessment in lymphoma," *J. Nucl. Med.*, 50(1):21-30, 2009.

[2] A. Rosset et al., "Informatics in radiology (infoRAD): navigating the fifth dimension: innovative interface for multidimensional multimodality image navigation," *Radiographics*, 26(1):299-308, 2006.

[3] S. Arens and G. Domik, "A survey of transfer functions suitable for volume rendering," *in* R. Westermann and G.L. Kindlmann (eds), *Volume Graphics*, pp. 77-83, 2010.

[4] B. Preim and D. Bartz, "Visualization in medicine: theory, algorithms, and applications," *Elsevier*, 2007.

[5] G. Kindlmann and J.W. Durkin, "Semi-automatic generation of transfer functions for direct volume rendering," *Proc. IEEE Symp. Volume visualization*, pp. 79-86, 1998.

[6] G. Kindlmann et al., "Curvature-based transfer functions for direct volume rendering: methods and applications," *Proc. IEEE Visualization*, pp. 513-20, 2003.

[7] C. Correa and K-L. Ma, "Size-based transfer functions: a new volume exploration technique," *IEEE Trans. Vis. Comput. Gr.,* 14(6):1380-7, 2008.

[8] C. Correa and K-L. Ma, "The occlusion spectrum for volume classification and visualization," *IEEE Trans. Vis. Comput. Gr.,* 15(6):1465-72, 2009.

[9] J.J. Caban and P. Rheingans, "Texture-based transfer functions for direct volume rendering," *IEEE Trans. Vis. Comput. Gr.,* 14(6):1364-71, 2008.

[10] W. Cai and G. Sakas, "Data intermixing and multi-volume rendering," *Computer Graphics Forum*, 18(3):359-68, 1999.

[11] R. Bramon et al., "Multimodal data fusion based on mutual information," *IEEE Trans. Vis. Comput. Gr., in press*, 2011.

[12] J. Kim et al., "Visualizing dual-modality rendered volumes using a dual-lookup table transfer function," *Comput.Sci.Eng.*, 9(1):20-5, 2007.

[13] U. Bordoloi and H.-W. Shen., "View selection for volume rendering," *IEEE Visualizatoin*, pp. 487-94, 2005.

[14] D. Ebert and P. Rheingans, "Volume illustration: non-photorealistic rendering of volume models", *Proc. IEEE Visualizatoin*, pp. 195-202, 2000.

[15] C. Correa and K-L. Ma, "Visibility histograms and visibility-driven transfer functions", *IEEE Trans. Vis. Comput. Gr.,* 17(2):192-204, 2011.

[16] J.D. Owens et al., "A survey of general-purpose computation on graphics hardware", *Comput. Graph. Forum*, 26(1):80-113, 2007.

[17] Voreen: volume rendering engine, www.voreen.org/

[18] S.C. Huang, "Anatomy of SUV. standardized uptake value", *J. Nucl. Med.*, 27(7):643-6, 2000.

[19] G.H. Weber et al., "Topology-controlled volume rendering", *IEEE Trans. Vis. Comput. Gr.,* 13(2):330-41, 2007.

[20] M.Ruiz et al., "Automatic transfer functions based on informational divergence," *EEE Trans. Vis. Comput. Gr.,* 17(12):1932-41, 2011.